

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
30 January 2003 (30.01.2003)

PCT

(10) International Publication Number  
**WO 03/009157 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 15/16**,  
15/173

(21) International Application Number: PCT/US02/22429

(22) International Filing Date: 15 July 2002 (15.07.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/305,992 16 July 2001 (16.07.2001) US  
60/305,969 16 July 2001 (16.07.2001) US  
10/000,708 31 October 2001 (31.10.2001) US  
10/000,709 31 October 2001 (31.10.2001) US

Pierce Street, Apartment 2101, Albany, CA 94706 (US).  
**MESSINGER, Adam**; 317 29th Street, #306, San Francisco, CA 94131 (US). **JACOBS, Dean, Bernard**; 1747 Madera Street, Berkeley, CA 94707 (US). **PULLARA, Sam**; 2030 3rd Street, #14, San Francisco, CA 94107 (US).

(74) Agents: **MEYER, Sheldon, R.** et al.; Fliesler Dubb Meyer and Lovejoy LLP, Suite 400, Four Embarcadero Center, San Francisco, CA 94111-4156 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.


(71) Applicant: **BEA SYSTEMS, INC.** [US/US]; 2315 North First Street, San Jose, CA 95131 (US).

(72) Inventors: **HALPERN, Eric, M.**; 160 Delmar Street, San Francisco, CA 94117 (US). **PEDDADA, Prasad**; 535

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK,

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR SESSION REPLICATION AND FAILOVER



(57) Abstract: A session replication system (100) provides real-time data replication without unnecessarily slowing down the user experience. A system in accordance with the present invention may utilize a primary server (104) to serve requests from a network client (102), as well as a secondary server (110) to replicate the session information. When a request is received on the session, an attempt may be made to serve the request on the primary server (104). If the primary is unable to receive or respond to the request, the request may be served on the secondary application server or on a new primary server. If the secondary server (110) receives the request, the secondary server may become the new primary server. If a new primary server is selected, the new primary may request the session information from the secondary server.

WO 03/009157 A1



TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**Published:**

— *with international search report*

## METHOD AND APPARATUS FOR SESSION REPLICATION AND FAILOVER

5

### CLAIM OF PRIORITY

[0001] This application claims priority to U.S. Provisional patent application No. 60/305,992, filed July 16, 2001, entitled METHOD AND APPARATUS FOR SERVLET SESSION REPLICATION AND FAILOVER; U.S. Patent application No. 10/000,708, filed October 31, 2001, entitled METHOD AND  
10 APPARATUS FOR SESSION REPLICATION AND FAILOVER; Provisional patent application No. 60/305,969, filed July 16, 2001, entitled HARDWARE LOAD-BALANCING APPARATUS FOR SERVLET SESSION REPLICATION; United States Patent Application No. 10/000,709, filed October 31, 2001 entitled HARDWARE LOAD-BALANCING APPARATUS  
15 FOR SESSION REPLICATION incorporated herein by reference.

### COPYRIGHT NOTICE

[0002] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no  
20 objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

25

### FIELD OF THE INVENTION

[0003] The invention relates generally to data replication and specifically to providing redundancy for a client network session.

### BACKGROUND

30

[0004] When a client connects to a server on a network and begins a session, there can be information stored on the server that is particular to

that client session. For example, a user of the client might place items in a virtual shopping cart. The selection of those items can be stored, at least temporarily, on the server. In this example, there is no need for other users or servers to have access to this information. It is desirable, however, that this data be highly available across a network or server cluster such that if the server storing the session data fails, it is possible to recover the data on another server.

**[0005]** One way to accomplish data recovery in such a situation is to store the information in a database during the session, although the information could also be stored by other means such as in a data file. Every time a change is made to the session data, an update is written to the database such that the data is accessible to every server having access to the database. The data is stored in a persistent place, and can easily be retrieved by another server.

**[0006]** A problem exists with this approach, however, in that it is fairly expensive to fetch session information from the database for each request. The multiple hits to the database can create a bottleneck and bog the system down to the point where it is basically inoperable, as the throughput of the system can depend on the number of database connections from the server. Also, these sessions may contain the type of information, to which users want quick access. With some applications, it is possible for there to be thousands of clients working simultaneously, resulting in thousands of concurrent sessions. Some servers are expected to host many different applications, which further increases the number of sessions that may need to be hosted.

**[0007]** It is desirable to improve the speed and efficiency of such a system so that these tens of thousands of users may use the system effectively. One way to avoid such a bottleneck is to assume that the servers will be up and running 99.9% of the time and simply neglect to backup any information. This may be the solution providing the fastest user

experience, but even 0.1% downtime resulting in data loss is unacceptable to many users.

#### BRIEF SUMMARY

5 [0008] Systems in accordance with the present invention can utilize a primary server to serve requests from a network client such as a web browser. The primary server can be selected from a pool of servers or server cluster. Once a primary server is selected, a client request can be served on that primary. A secondary session server can then be chosen,  
10 such as by the primary server. Once the primary server responds to the request, the information relating to the session is sent from the primary server to the secondary session server. This can be a full set of information for a first request on a session, or can simply be an update to existing information in a session in response to subsequent requests.  
15 Information identifying the primary and secondary servers can be stored on the client, such as a "token" that is stored as a cookie or passed on top of standard RMI in a manner similar to a transaction or security context. This identification information or "token" can accompany each request.

[0009] A system can take advantage of load balancing using either  
20 hardware or software. In a process useful with software load balancing, a request may be received on a session for which a primary and secondary session server have already been selected. An attempt is made to serve the request on the primary server. If the primary is unable to receive or respond to the request, the request can be served on the secondary  
25 application server. If the secondary server receives the request, the secondary server becomes the new primary server. A new secondary server can be selected and sent the session information from the new primary server in order to maintain redundancy.

[0010] In a process useful with a hardware load balancer, a request is  
30 received on a session for which a primary and secondary session server have been selected. An attempt is then made to serve the request on the

primary server. If the primary server is unable to receive or respond to the request, the hardware load balancer can select a new primary server and attempt to serve the request on the new primary server, instead of using the secondary session server. The session information can be sent from the secondary session server to the new primary server, such as in response to a request from the new primary. The new primary server can then respond to the request, and send updated session information to the secondary server, so that the servers are in sync with regard to that session.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIGURE 1 is a diagram of an application server system in accordance with one embodiment of the present invention.

[0012] FIGURE 2 is a diagram of a multi-level architecture in accordance with one embodiment of the present invention.

[0013] FIGURE 3 is a diagram of a servlet engine system in accordance with one embodiment of the present invention.

[0014] FIGURE 4 is a diagram of a load balancer system in accordance with one embodiment of the present invention.

[0015] FIGURE 5 is a diagram of a Java system in accordance with one embodiment of the present invention.

[0016] FIGURE 6 is a flowchart for a process in accordance with one embodiment of the present invention.

[0017] FIGURE 7 is a flowchart for a software load balancer process in accordance with one embodiment of the present invention.

[0018] FIGURE 8 is a flowchart for a hardware load balancer process in accordance with one embodiment of the present invention.

#### DETAILED DESCRIPTION

[0019] The present invention overcomes many of the deficiencies of prior art replication systems. In one system in accordance with an embodiment

of the present invention, a session is created when a client makes a request to a server on a network, such as a local area network (LAN), ethernet, or Internet. The session server receiving the request can consist of any server that may be used to store information for a session and/or  
5 generate a response to a session request, such as for example an application server, a web server, an object server, or a servlet engine. The server ultimately receiving the request becomes the "primary" server, or the server to which the client will send future requests. The system can then choose a "secondary" server for that session, which acts as a source for  
10 redundancy.

[0020] Each time an update is made in that session, the change may not only be stored on the primary server, but can also be sent, such as by remote call, to the secondary server. The entire set of session data need not be forwarded to the secondary server each time a change is made, but  
15 only that data or information that has changed, such as can be sent in a delta or packet of information. Sending the minimum amount of necessary information in the delta can improve overall system efficiency. The replication acts like a mirroring system, except for the fact that it acts on session data. This mirroring can be done, in one example, for web  
20 applications using a servlet engine.

[0021] When a client connects to a server, a session object can be created that is associated with the client or user. The session object can be maintained on the primary server for the duration of the session or can be timed out after a specified period of time. Each session object can be  
25 given a unique identifier or identification number to identify the client and/or object to the server. The server chosen to serve the request can act as the primary server for the duration of the session. The primary server can choose a secondary server for the session object, such that each time the object is updated the update is also stored on the secondary server. The  
30 secondary server can be optimized to receive only minimum information or to batch updates, in order to improve the efficiency of the system.

[0022] One web-based system 100 in accordance with an embodiment of the present invention is shown in **Figure 1**. In this system, a browser 102, or client, makes a request that is received by a web server 104. The web server 104 acts as a proxy, in that the web server looks at the request and determines which of the object servers 110 should receive the request. The web server can have a plug-in, or plug-in API, that is aware of the request. A plug-in is generally an object that is added to an application in order to provide additional functionality without the need to launching any external applications. The plug-in can make a load-balancing decision in order to choose between the object servers 110 which are available to create and house the session for the client 102. The web server 104 proxies back to the chosen object server 110, which can be housed in an application server 106. A servlet engine 108 in the application server 106 can execute servlets that invoke objects on the object server 110 in order to respond to the request. In order to fully respond to the request, the object server 110 may also need to pull information from a database 112 or data store. The object server 110 can create the session upon receiving the request. In order to provide security, the application server 106 and database can be located behind a firewall 114, as is known and used in the art.

[0023] The object server in this example then chooses a secondary server for the session. In an alternative embodiment, the plug-in can be used to choose a secondary server. The plug-in can also use load balancing for the decision.

[0024] The object server passes the data to the secondary server, and lets the secondary server know it will be the backup. The object server then creates a cookie to be sent to, and stored by, the client. The cookie contains an identification of the primary and secondary servers used for the session.



**[0025]** When a subsequent request is sent by the client on the same session, it does not matter which web server receives the request. The web server will look at the cookie to determine the primary server for that session, and will then route the request to that primary server.

5 **[0026]** Assume an example, as shown in **Figure 3**, with three servlet engines **306**, **308**, and **312**, each being capable of acting as a primary server. If the session is current on primary server **306**, but primary server **306** fails, the web server **304** can determine which server was chosen as the secondary server by examining the cookie information sent with the  
10 request from the browser **302**. The web server can then try to route the request to secondary server **308**, which also contains the session state information **310**. The web server can return a response to the browser **302**, which will make another request that can be directed by the web server **304** to the secondary server **308**. If secondary server **308** receives  
15 the request, it can become the new primary server automatically, as secondary server **308** knows it only receives a request directly from the web server if the primary server **306** fails to accept the request. At this point, secondary server **308**, now the primary server **308**, can choose new secondary server **312**. Alternatively, the plug-in to the web server **304** can  
20 choose new secondary server **312**. One possible place for a communication breakdown is shown by a first virtual boundary **314**, which exists between the browser/client **302** and the web servers **304**. A second virtual boundary exists between the web servers **304** and the servlet engines **306**, **308**, and **312**.

25 **[0027]** In some embodiments, the secondary server or web server actively monitors the primary server in order to determine the status of the primary server. This monitoring can be done by any appropriate manner, such as by continually or periodically "pinging" the primary server to determine whether it is connected to the network. If it is determined that the primary  
30 server is unable to accept requests, the secondary server can become the

new primary server. A new secondary server can then be selected. One advantage of such a design is that the window of time in which a dual server failure could result in session state loss is narrowed. While some embodiments allow the window to be defined by the client request rate, this approach would allow the window to be defined by the rate of server  
5 pinging.

**[0028]** The new primary and secondary servers are similarly responsible for information pertaining to the session. The server that was previously the primary server may no longer have any responsibility or information for that  
10 session, even if that server becomes able to accept and process requests while the session is still current. The secondary server may automatically change its state, such that it becomes the new primary server for the session, but it may choose not to assign a new secondary server until the new primary server receives a request.

**[0029]** It may be undesirable to actively create a new secondary server or backup the session information on a secondary server, as it may not be known whether the new primary server will receive another request. Creating a new secondary server or backing up information that will not get  
15 used can unnecessarily waste resources. The session may alternatively be short lived, and may not "live" long enough to receive a subsequent request. Each session typically has a time-out value, such that if the session is inactive for a specified period of time it will "time-out" or "die", the session will be ended, and all data stored for the session may be deleted to conserve memory. In such a case, not only might the creation of a  
20 secondary server waste resources, but it may also require unnecessary "clean-up" work to remove the session information from the new secondary server.

**[0030]** The primary and/or secondary server can be chosen according to an algorithm, which may have as options, for example, any server in a  
30 specified server cluster. While it may be efficient for an algorithm to choose the primary and secondary server for each session, there are

cases where the input of an administrator may be desirable. For instance, it may be possible that multiple servers are located on one machine. If an algorithm is choosing servers, such as an algorithm based on load, the algorithm may select two servers on the same machine. In the event of a machine failure, both servers may be unavailable and the session data may be unavailable and/or lost. An administrator, however, can choose to specify primary and secondary servers that are on different machines. This can provide for redundancy not only across servers, but across machines as well.

10 **[0031]** Alternatively, it may be possible to build a parameter in the algorithm itself that, when doing a load-balancing analysis, takes into account the machine on which a server is located. If the server with the lowest current load is on the same machine as the primary server, the algorithm can choose to go to the server with the lowest load that is on a different machine. This approach can be expanded to any level of separation, such as servers in different rooms, different buildings, or different cities.

15 **[0032]** In order to allow servers in a cluster to function independently, the servers can be loosely coupled. To achieve this loose coupling, each server in the cluster can be configured to detect the status of other cluster servers such that action can be taken when a server leaves the cluster, either voluntarily or involuntarily. In one embodiment, servers can rely on the underlying operating system to monitor the status of the cluster servers. Other embodiments can require the servers to do monitoring. Embodiments where cluster servers do not have to participate in cluster monitoring may be preferred, as the servers' resources are available to improve the overall throughput of the system.

20 **[0033]** **Figure 2** shows a multi-tier cluster architecture **200** in accordance with the present invention. Each object in the system can be clustered by making instances of the object available on several servers. The architecture is shown to include virtual boundaries. The term "virtual boundary" refers to a place where a network connection may fail.

25

30

**[0034]** In Figure 2, the first virtual boundary 212 is shown between the browsers 202 and the web servers 204. A second barrier 214 is shown between the web servers 204 and the servlet engines 206. A third barrier 216 is shown between the servlet engines 206 and the object servers 208.

5 Finally, a fourth barrier 218 is shown between the object servers 208 and the database 210. Each barrier indicates a possible point for communication failure, that may also be able to take advantage of load balancing.

**[0035]** At the first virtual barrier, it is possible that a browser may not be able to get to a specific web server. This may not be a problem in a system in accordance with the present invention, however, because the information relating to the primary and secondary servers may already be stored in a cookie in the browser. The browser can contact any web server on the network, because the browser can indicate to the web server, through the cookie, which server should receive the request. While the system may be most efficient on a local area network (LAN), a similar approach can be used on any capable network. For instance, it may be possible for the browser to contact a second web server and/or backend servers over the Internet that might be located in separate buildings from the first web server.

**[0036]** Depending on the application, the primary and secondary servers can be of several different server types, such as web servers, servlet engines, or Enterprise Java bean ("ejb") engines. It may still be possible for each server in a cluster to be separate and specialized, such as being of different server types, but still be capable of acting as a primary and/or secondary server.

**[0037]** If clustering is enabled on a system in accordance with the present invention, it may be possible to transparently add new servers to the system to act as additional primary and secondary servers. Clustering is, generally, an approach to server management that allows management of a set of servers by establishing a "managing" server in that set of servers.

This approach can simplify the deployment and synchronization of potentially diverse components among the servers in the cluster. Clustering can substantially improve system reliability and scalability.

5      **[0038]** When clustering with a system in accordance with the present invention, each server in a cluster can be configured to detect a new server entering the cluster and designate that new server as a secondary server to any existing primary server. The method used for load balancing may can immediately designate the new server as a primary or secondary server.

10      **[0039]** Systems in accordance with the present invention may alternatively utilize a hardware load balancer to direct incoming requests. In an Internet setting, for example, a hardware load balancer can sit on the network with an IP (Internet Protocol) address. Incoming requests from browsers or clients can be directed to that IP address. The hardware load balancer can  
15      then redirect those requests to other IP addresses, or other servers each assigned an IP address, located in the system but "behind" the hardware load balancer. In this way, it appears to the browser as if the request is always going to the same IP address, when in fact it may be going to multiple servers behind that IP address. The hardware load balancer can  
20      be aware of all servers located behind it in the network, such as may be the result of hardwiring the servers to the hardware load balancer, instead of utilizing another method such as software clustering.

25      **[0040]** There may be advantages to using a hardware load balancer. A hardware load balancer can utilize better algorithms for load balancing than other approaches. The hardware load balancer may be able to detect node failures, so that those nodes may be pulled out of the list of servers available to the algorithm. This node removal can prevent the algorithm from trying to go to servers that may not be reachable, even though those particular servers might not yet have been sent a request.

30      **[0041]** A system in accordance with the present invention can also use a Domain Name System (DNS) protocol, such as DNS Round Robin, instead

of using a hardware load balancer to map a domain name to several IP addresses, or to redirect requests sent to a web server to several object servers. DNS, however, does not typically determine or detect whether those IP addresses are actually "live".

5     **[0042]** A hardware load balancer can be used to proxy certain types of requests to specific servers or server clusters, depending on whether the request requires dynamic page generation or whether the request is for a static page. In **Figure 4**, a load balancer **414** is shown between the web browser **402** and the web servers **404**, **408**, **412**.

10    **[0043]** While it may be desirable to optimize a hardware load balancer **414** for use with the present invention, it may be undesirable to require physical changes to the load balancer itself. It may also be undesirable for a hardware load balancer to have to read cookies and figure out that if the first primary **404** fails, the request needs to be redirected to the secondary  
15    server **408** indicated in the cookie stored on the browser **402**. It may, however, be desirable to have the load balancer direct the request where the load balancer wants, and then make sure the system recovers appropriately.

20    **[0044]** In one such approach, the hardware load balancer **414** tends to send requests from a browser **402** or client to one server, based on some arbitrary information stored in the cookie on the web browser. For example, a cookie can have an initial string of information, followed by a segment of information related to the primary and secondary servers, as well as a session identifier used for replication. The hardware load  
25    balancer **414** can be configured to look only at this segment of information. If this segment of information does not change between successive cookies, the load balancer may keep redirecting the requests back to the primary server **404**. Such "session stickiness" can also be based on other appropriate schemes, such as may utilize the IP address of the client.

30    **[0045]** The segment of information in the cookie can remain the same as long as the requests can go back to the primary server. If the primary

server fails for any reason, the second server can assign itself to be the new primary server. The new primary server can then insert new information in the segment for a new secondary server, which can be selected by the new primary server or load-balancing mechanism.

5 Alternatively, the hardware load balancer can choose a new primary server and redirect the request to the new primary server.

[0046] The first request of a session that is received by the load balancer may not be hard coded to go to one server primarily. The decision to “stick” to one server can be made after the first request is made and comes  
10 back from the object server or other backend server. Hardware load balancers can be smart enough to do this “simple stickiness”, or to return primarily to the assigned primary server with which the load balancer has a connection.

[0047] If no cookie exists, the hardware load balancer can be configured  
15 to use any of a number of load balancing methods, such as may be based on load or response time. The load balancer can then select a server, such as in the appropriate cluster, and direct the request to that server. When that primary server responds to the request, the server can send the browser a cookie containing the segment of information relating to the  
20 primary and secondary servers. Each subsequent request from that browser that is received by the hardware load balancer can have that cookie associated with it, such that the load balancer can associate that request with the primary server.

[0048] The system still may not be able to guarantee that a request will go  
25 to the primary server. As shown in **Figure 4**, if a failure occurs at the primary server **404** and another request comes into the load balancer **414**, the load balancer can simply make another load balancing decision and direct the request to another server **412**. The request may not go to the second server **408**. This approach is different from that described above  
30 for a plug-in approach, where the request can automatically go to the second server. In this way, a hardware load balancer is less “intelligent”

than a special proxy plug-in, similar to those described above.

[0049] If the server 412 chosen by the load balancer 414 is not the second server 408, the chosen server 412 may realize that the request is a request on a session that it is not hosting. In this case, the chosen server 412 can  
5 look to the cookie in order to determine the secondary server 408.

[0050] Once the chosen server 412 has located the secondary server 408, the chosen server 412 can request session state information 410 from the secondary server 408. The chosen server 412 can then transform itself into the new primary server for the session. The secondary server 408 in  
10 this case can remain the same. The cookie is updated so the load balancer 414 will keep directing the requests to the new primary server 412.

[0051] In the event that the load balancer 414 chooses to direct the request to a new server that happens to be the secondary server 408, the  
15 secondary server can set itself as the new primary server and a new secondary server can be selected.

[0052] A system with a hardware load balancer having a servlet cluster behind it can provide a fast data path. If a web server does the routing, it may be necessary for the request to come up into the software where  
20 some code is executed, and then be sent back out onto the network. The load balancer/servlet cluster system does everything at a low, protocol level so it may be comparatively very fast.

[0053] It may be advantageous to have the load balancing algorithms as localized as possible. In the hardware load balancer case, it may only be  
25 necessary to ensure that the software on the server is operating properly, such as software that may be written in Java. In systems without a hardware load balancer, it may be necessary to make sure that each special plug-in of every web server in the system works as well.

[0054] It may also be necessary to support plug-ins for different platforms.  
30 A hardware load balancer can work equally well with systems based on



differing platforms, such as a Netscape Application Server (NAS), WebLogic Server™ (WLS), Microsoft® Internet Information Server (IIS), or Apache HTTP Server. With the hardware load balancer, the system can be reduced by one level of complexity, as one of the banks may be removed. This is shown in Figure 4, where the web server and servlet engine are in the same process.

**[0055]** Some of the systems described above can utilize servlets for web access. A similar mechanism can be used for accessing stateful session beans, a type of Enterprise Java bean ("ejb"). While servlets can be used to service requests from browser clients, ejb servers can be used to support requests from Java clients.

**[0056]** With Java clients, there can be a single, persistent connection for the entire duration of a session. There may then be no need (or support) for cookies. Also, since a persistent connection exists, there may no longer be a need for a load balancer. A Java client can connect to one of the backend servers using, for example, DNS or a load balancer. The Java client can then look up a "handle" to a stateful session. A handle in Java is similar to a pointer, which may be used to locate the appropriate session.

**[0057]** Referring to the system 500 of Figure 5, once a Java client 502 is connected to a handle, the stateful session bean 510 can be created. The stateful session bean 510 can be used to handle the caching or storing of information for the session. When the stateful session bean is created, the server housing the bean can become the primary server 508, to which the Java client 502 can make requests. The primary server 506 can then choose a secondary server 512. The secondary server 510 can also have a stateful session bean 512 to cache or store the session information.

**[0058]** The stateful session bean 508 can pass this information back to the Java client 502, similar to sending a cookie, using an RMI (Java Remote Method Invocation) protocol. Extra information can be placed "on top" of the standard RMI in order to make this cookie simulation work, similar to

the way transaction context propagation works. The primary/secondary server identifier pair can be passed back to the Java client 502 with every response. Each time the Java client 502 makes a call, it can call through an interface 504 into special RMI code adapted to continue to make calls to the primary server 506 for that session. If the primary fails, the Java client 502 can look at the information regarding the location of the secondary server 510, and can instead make a request to the secondary server. It may be preferable, for efficiency, if only the essential information regarding server identification is passed back on top of the RMI.

[0059] The Java client 502 may always know which server is the secondary server 510. The Java client can have much of the same logic that a proxy might have, such as always knowing to go to the secondary server 510 if the primary server 506 is unavailable. The Java client can monitor server health in order to avoid sending requests to unavailable servers. In the event that the secondary server 510 becomes the new primary server, a new secondary server 514 can still be chosen. The logic for selecting a new primary and/or secondary server can be similar to that described above. The Java client can immediately update to the new primary/secondary servers.

[0060] In accordance with the above discussion, systems in accordance with the present invention can generally follow one of two branching paths, although variations including those mentioned above may be utilized. A common part of such a path is shown in **Figure 6**. In the process 600 of **Figure 6**, a primary server is selected from a group of servers 602. Once a primary server is selected, the client request is served on that primary server 604. A secondary server is then chosen 606, possibly by the primary server. The session information is then sent from the primary to the secondary server 608. Information identifying the primary and secondary servers can be stored on the client 610, such as may be stored in a cookie or passed on top of standard (or other) RMI.

[0061] From this point, the process branches off into a path that can be useful for software load balancing, and a process that can be useful for hardware balancing. **Figure 7** shows a process **700** useful for software load balancing. In the process **700**, a request is received on a session for which a primary and secondary server have already been selected, and the identification of the primary is garnered from the information stored on the client **702**. An attempt is then made to serve the request on the primary server **704**. If the primary is unable to serve the request, the request is served on the secondary server **706**. Once the secondary server receives the request, the secondary server becomes the new primary server **708**. A new secondary server is then selected and sent the session information from the new primary server **710**.

[0062] Another path, useful for systems with hardware load balancers, is shown in **Figure 8**. In the process **800** of **Figure 8**, a request is received on a session for which a primary and secondary server have already been selected, and the identification of the primary server is garnered from the information stored on the client **802**. An attempt is then made to serve the request on the primary server **804**. If the primary server is unable to serve the request, the hardware load balancer selects a new primary server and attempts to server the request on the new primary server **806**. The session information is then sent from the secondary server to the new primary server **808**, such as in response to a request from the new primary. The new primary server may then respond to the request, and send updated session information to the secondary server **810**.

[0063] To maintain the consistency in any embodiment of the present invention, a change in session data can be associated with a version number. The primary and secondary servers may each know which version of the session it is storing. A server can be instructed to modify the data only if it receives a request that has a version number later, or higher, than the one it is currently storing. The primary and secondary servers can

periodically check each other to ensure that they are both on the same version number. The version number can use a method as simple as incrementing a number to guarantee ordering. For session information to remain consistent, it may be desirable for the primary and secondary servers to be in synchronization. When a version number is out of synchronization, the primary server can choose to send the entire session information to the secondary server to bring the session back into synchronization. The synchronization also facilitates the ability of the servers to switch roles between primary and secondary if the need arises.

5 [0064] If the primary server is unable to update information on the secondary server, for reasons such as a bad connection, it may be possible that the primary server will keep updating and the secondary server will be unaware of any updates. It may then be possible for the primary server to be several versions ahead of the secondary server.

10 [0065] Once the primary server is again able to send information to the secondary server, a delta between two successive versions may not work. In such a case, the primary server can send an entirely new set of session data to the secondary server in order to make the data session consistent across both servers. In this case, the secondary server either gets a delta

15 between successive versions, or it gets all the data for the entire session. In other embodiments, it may be possible to generate a delta between arbitrary versions in order to bring the secondary server up to the current version.

20 [0066] In simulating cookies for tracking Java states, large random numbers can be used for server identification. The numbers can be large enough that it is highly unlikely that the sum of two different pairs of identification numbers will be the same. It may be possible to only send these two numbers back to the Java client, and the server pair can be identified by adding the two numbers together to get a new number. This

25 allows the passing of only one number to identify two servers, which can improve efficiency. Since the Java client can have a persistent connection

30

with a specific server, the client can identify the second server by subtracting the identification number of the primary server from the summed number being passed, in order to arrive at the identification number of the second server.

5     **[0066]** Java objects such as session beans may, however, be stateless or have a transient state, as opposed to the persistent state discussed above. If a Java session bean is stateless, the bean may not be able to maintain session information between invokes, or successive requests. If the session information is stored elsewhere, the stateless beans can load the  
10     session information temporarily in order to serve the request. Failover, or turning session control over to a new primary server having replicated session information, can only occur where there was a clean invoke failure, such as where the primary server never received the request, the request was transactional and was aborted, or was a one-time-only request. If the  
15     session bean is transient, on the other hand, instances can be created by a stateless factory with stateless load balancing and failover. The bean in a transient state may either not be backed up or may be backed up in-memory using primary/secondary replication as discussed above.

20     **[0067]** Batch updates can be used to improve the throughput of the system with an increased failure window. When batching, or "boxcarring," several requests are sent together as one large request in order to improve efficiency and scalability. The batching of requests can be based upon any of a number of criteria, such as time intervals or numbers of requests. For instance, a system can send a batch of requests every 10 seconds or for  
25     every 100 individual session update messages. The system can also accommodate both criteria, sending a batch when either 10 seconds has passed since the last batch or when 100 requests are received, whichever comes first. Batching may cause the system to no longer be as reliable as synchronous updates, but can improve the overall system scalability.

30     **[0068]** The criteria may also be configurable, such as by a user or an administrator. Configurable criteria can be appropriate in situations where

a system encounters a lot of traffic at certain times, but little to no traffic at other times. Configurable criteria can, for example, allow batching of every 100 messages at peak time, but no batching at all during off time, so that every request is sent in a reasonable amount of time.

5     **[0069]** A system administrator can also choose to pair two servers in a cluster as primary and secondary servers. The input of the administrator can be desirable in order to improve the overall fault tolerance of the system. For example, multiple servers can be located on one physical machine and an algorithm might choose to place both primary and  
10     secondary servers on the same machine. The session information could then be lost entirely if that machine fails. In order to prevent the loss of session information from machine failures, an administrator can choose to specify a primary server and a secondary server each on a separate physical machine. The administrator can also choose a primary server  
15     based on various load-balancing schemes. Examples of possible schemes are based on server load, connection number, and physical proximity.

**[0070]** The foregoing description of preferred embodiments of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise  
20     forms disclosed. Obviously, many modifications and variations will be apparent to the practitioner skilled in the art. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various  
25     modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalence.

CLAIMS

- 1     1.     A system for replicating information in a client session, comprising:  
2             b.     a primary server adapted to receive a request in a client  
3                     session and serve a response to that request, said primary  
4                     server further adapted to store session information for that  
5                     client session;  
6             c.     a secondary server adapted to receive request in a client  
7                     session and serve a response to that request, said  
8                     secondary server further adapted to store session  
9                     information for that client session; and  
10            d.     a web server adapted to receive a request from a client, the  
11                    request containing identification information for said primary  
12                    and secondary servers, the web server adapted to process  
13                    the identification information and serve the process request  
14                    on said primary server, the web server further adapted to  
15                    serve the request on said secondary server if said primary  
16                    server is unable to process the request.
- 1     2.     A system according to claim 1, further comprising a database in  
2     communication with said primary and secondary servers, said database  
3     storing information useful in processing the request.
- 1     3.     A system according to claim 1, wherein said primary server is further  
2     adapted to update the session information stored in said second session  
3     server each time said primary server receives a request for that client  
4     session.
- 1     4.     A system according to claim 1, wherein said web server is further  
2     adapted to choose said primary server from a plurality of servers when an

3 initial request is received for the client session and no identification  
4 information exists on the client.

1 5. A system according to claim 1, further comprising a cookie adapted  
2 to be stored on the client, said cookie adapted to contain the identification  
3 information for said primary and secondary servers.

1 6. A system according to claim 5, wherein said primary server is further  
2 adapted to generate said cookie on the client.

1 7. A system according to claim 1, wherein said web server further  
2 comprises a plug-in containing an algorithm to be used in choosing said  
3 primary server.

1 8. A system according to claim 1, wherein said web server further  
2 comprises a plug-in containing an algorithm to be used in directing the  
3 request to said primary and secondary servers.

1 9. A system according to claim 7, wherein said algorithm is a load-  
2 balancing algorithm.

1 10. A system according to claim 1, wherein said primary server is  
2 adapted to select said secondary server.

1 11. A system according to claim 1, wherein said primary server is  
2 adapted to start the client session when receiving an initial request from  
3 the client.

1 12. A system according to claim 1, wherein said secondary server is  
2 further adapted to become the new primary server if the secondary server  
3 receives a request that could not be processed by said primary server.



- 1     13.   A system according to claim 12, wherein said secondary server is  
2     further adapted to select a new secondary server and send the session  
3     information to the new secondary server.
- 1     14.   A system according to claim 12, wherein said web server chooses  
2     a new secondary server.
- 1     15.   A system according to claim 1, wherein said secondary server is  
2     further adapted to monitor said primary server to determine if said primary  
3     server may receive requests.
- 1     16.   A system according to claim 1, wherein said secondary server is  
2     further adapted become a new primary server if said primary server is  
3     unable to receive requests.
- 1     17.   A system according to claim 1, wherein one of said primary and  
2     secondary servers is selected from the group consisting of web servers,  
3     servlet engines, and enterprise Java bean engines.
- 1     18.   A system for replicating information in a Java client session,  
2     comprising:  
3         a.   a primary server adapted to receive a request from a Java  
4             client in a client session and serve a response to that  
5             request, said primary server further adapted to store session  
6             information for that client session;  
7         b.   a secondary server adapted to receive request in a client  
8             session and serve a response to that request, said primary  
9             server further adapted to store session information for that  
10            client session; and.

11           c.     a stateful session bean on the primary server used to store  
12                   information for the client session and pass the identity of said  
13                   primary and secondary servers back to the client.

1     19.   A system according to claim 18, further comprising a load balancer  
2     adapted to receive a request from a client and select said primary server.

1     20.   A system according to claim 18, wherein said primary server is  
2     adapted to maintain a persistent connection with the Java client during the  
3     client session.

1     21.   A system according to claim 1, wherein said web server is further  
2     adapted to send requests to said primary server in batches.

1     22.   A method for providing redundancy in a client session, comprising:  
2           a.     making a load balancing decision for an initial request from  
3                   a client in a client session in order to select a primary server  
4                   from a plurality of session servers;  
5           b.     serving said request on said primary server;  
6           c.     selecting a secondary server;  
7           d.     sending session information for that client session from the  
8                   primary server to the secondary server; and  
9           e.     updating session information on the primary and secondary  
10                  servers each time a request is received in the client session.

1     23.   A method according to claim 22, further comprising:  
2     storing the information in a cookie on the client.

1     24.   A method according to claim 22, further comprising:  
2     serving the request on the secondary server if the primary server is  
3     unable to process the request.

- 1     25.    A method according to claim 24, further comprising:  
2            selecting a new secondary server.
- 1     26.    A method according to claim 22, further comprising:  
2            serving requests to the primary server in batches.
- 1     27.    A method according to claim 22, further comprising:  
2            sending session information for that client session from the primary  
3            server to the secondary server in batches.
- 1     28.    A method according to claim 22, further comprising:  
2            storing session information in a stateful session bean.
- 1     29.    A method according to claim 22, further comprising:  
2            associating a version number with each update to the session  
3            information in response to a request.
- 1     30.    A method according to claim 22, wherein sending session  
2            information for that client session from the primary server to the secondary  
3            server comprises sending a delta of information containing the changes in  
4            the session information.
- 1     31.    A method according to claim 22, further comprising:  
2            assigning the primary and secondary servers each an identification  
3            number.
- 1     32.    A method according to claim 31, further comprising:  
2            adding the identification number for the primary and secondary  
3            servers to obtain one number that describes both the primary and  
4            secondary servers.

- 1     33.   A system for replicating information in a client session, comprising:
- 2           a.     a plurality of session servers;
- 3           b.     a primary server in said plurality of session servers, said
- 4                 primary server adapted to receive a request in a client
- 5                 session and serve a response to that request, said primary
- 6                 server further adapted to store session information for that
- 7                 client session, and adapted to generate a cookie containing
- 8                 information relating to said primary server and transmit the
- 9                 cookie to the session client;
- 10          c.     a secondary server in said plurality of session servers
- 11                 selected by said primary server, said secondary server
- 12                 adapted to receive the session information from said primary
- 13                 server and store the session information, the secondary
- 14                 server further adapted to receive a request in a client session
- 15                 and serve a response to that request; and
- 16          d.     a web server adapted to receive a request from a client, the
- 17                 request containing identification information for said primary
- 18                 and secondary servers, the web server adapted to process
- 19                 the identification information and serve the process request
- 20                 on said primary server, the web server further adapted to
- 21                 serve the request on said secondary server if said primary
- 22                 server is unable to process the request.
- 1     34.   A method for providing redundancy in a client session, comprising:
- 2           a.     making a load balancing decision for an initial request from
- 3                 a client in a client session in order to select a primary server
- 4                 from a plurality of session servers;
- 5           b.     serving said request on said primary server;

- 6           c.     sending session information for that client session from the
- 7                 primary server to a secondary server selected by said
- 8                 primary server
- 9           d.     storing a cookie on the session client containing information
- 10                identifying the primary and secondary servers; and
- 11           e.     updating session information on the primary and secondary
- 12                servers each time a request is received in the client session.

1     35.    A method according to claim 34, further comprising the step of  
2     reading identification information in a request received in the client session  
3     and serving the request on the primary server.

- 1     36.    A system for replicating information in a client session, comprising:
- 2           a.     a plurality of session servers;
  - 3           b.     a primary server in said plurality of session servers, said
  - 4                 primary server adapted to receive a request in a client
  - 5                 session and serve a response to that request, said primary
  - 6                 server further adapted to store session information for that
  - 7                 client session, and adapted to generate a cookie containing
  - 8                 information relating to said primary server and transmit the
  - 9                 cookie to the session client;
  - 10          c.     a secondary server in said plurality of session servers
  - 11                selected by said primary server, said secondary server
  - 12                adapted to receive the session information from said primary
  - 13                server and store the session information, the secondary
  - 14                server further adapted to receive a request in a client session
  - 15                and serve a response to that request; and
  - 16          d.     a web server containing load-balancing logic for selecting
  - 17                said primary server, said web server adapted to receive an
  - 18                initial request from a client and direct that request to said
  - 19                primary server, said web server further adapted to receive a

20 subsequent request from a client, the subsequent request  
21 containing identification information for said primary and  
22 secondary servers, the web server adapted to process the  
23 identification information and serve the process request on  
24 said primary server, the web server further adapted to serve  
25 the request on said secondary server if said primary server  
26 is unable to process the request.

1 37. A method for providing redundancy in a client session, comprising:  
2 a. making a load balancing decision for an initial request from  
3 a client in a client session in order to select a primary server  
4 from a plurality of session servers;  
5 b. serving said request on said primary server;  
6 c. sending session information for that client session from the  
7 primary server to a secondary server selected by said  
8 primary server;  
9 d. storing a cookie on the session client, the cookie containing  
10 information identifying the primary and secondary servers;  
11 e. reading the identification information received with any  
12 subsequent request in the client session and serving that  
13 subsequent request on the primary server; and  
14 f. updating session information on the primary and secondary  
15 servers each time a subsequent request is served on the  
16 primary server for that client session.

1 38. A method for providing redundancy in a client session, comprising:  
2 a. selecting a primary server from a plurality of session servers  
3 in response to an initial request from a client;  
4 b. serving said request on said primary server to begin a client  
5 session;  
6 c. selecting a secondary server;

- 7 d. sending session information for that client session from the  
8 primary server to the secondary server; and  
9 e. updating session information on the primary and secondary  
10 servers each time a request is received in the client session.

1 39. A method according to claim 38, further comprising:  
2 storing the information in a cookie on the client.

1 40. A method according to claim 38, further comprising:  
2 serving the request on the secondary server if the primary server is  
3 unable to process the request.

1 41. A method according to claim 40, further comprising:  
2 selecting a new secondary server.

1 42. A method according to claim 38, further comprising:  
2 serving requests to the primary server in batches.

1 43. A method according to claim 38, further comprising:  
2 sending session information for that client session from the primary  
3 server to the secondary server in batches.

1 44. A system for replicating information in a client session, comprising:  
2 a. a plurality of servers;  
3 b. a primary server in said plurality of servers, said primary  
4 server adapted to receive a request in a client session and  
5 serve a response to that request, said primary server further  
6 adapted to store session information for that client session;  
7 c. a secondary server in said plurality of servers, said  
8 secondary server adapted to receive request in a client

9 session and serve a response to that request, said primary  
10 server further adapted to store session information for that  
11 client session; and  
12 d. a hardware load balancer adapted to receive a request from  
13 a client, the request containing identification information for  
14 said primary and secondary servers, said hardware load  
15 balancer adapted to check that portion of the request  
16 containing the identification information and serve the  
17 process request on said primary server if that portion has not  
18 changed since a previous request, said hardware load  
19 balancer further adapted to select a new primary server from  
20 said plurality of servers if that portion has changed.

1 45. A system according to claim 44, further comprising a cookie adapted  
2 to be stored on the client, said cookie adapted to contain the identification  
3 information for said primary and secondary servers.

1 46. A system according to claim 45, wherein said cookie contains a  
2 number for said primary server and a number for said secondary server.

1 47. A system according to claim 45, wherein said cookie contains a  
2 number that is the sum of a number for said primary server and a number  
3 for said secondary server.

1 48. A system according to claim 44, wherein said hardware load  
2 balancer is further adapted to select a new primary server from said  
3 plurality of servers if said primary server cannot process the request.

1 49. A system according to claim 44, wherein said primary server is  
2 adapted to request session information for that client session from said  
3 secondary server when said primary server receives a request on a client



4 session that the primary server is not hosting, information for that client  
5 session being stored on said secondary server.

1 50. A system according to claim 44, wherein said primary server is  
2 further adapted to read a cookie associated with a request received for a  
3 client session and determine whether the primary server is hosting that  
4 client session.

1 51. A system according to claim 44, wherein said hardware load  
2 balancer is further adapted to send requests to said primary server in  
3 batches.

1 52. A method for providing redundancy in a client session, comprising:  
2 a. making a load balancing decision for an initial request from  
3 a client in a client session in order to select a primary server  
4 from a plurality of servers, the load balancing decision being  
5 made using an algorithm in a hardware load balancer;  
6 b. serving said request on the primary server;  
7 c. selecting a secondary server using the primary server;  
8 d. sending session information for that client session from the  
9 primary server to the secondary server; and  
10 e. updating session information on the primary and secondary  
11 servers each time a request is received in the client session.

1 53. A method according to claim 52, further comprising:  
2 storing the information in a cookie on the client.

1 54. A method according to claim 52, further comprising:  
2 selecting a new primary server using the hardware load balancer if  
3 the primary server is unable to process the request.

- 1 55. A method according to claim 54, further comprising:  
2 serving the request on the new primary server.
- 1 56. A method according to claim 54, further comprising:  
2 selecting a new secondary server.
- 1 57. A method according to claim 52, further comprising:  
2 serving requests to the primary server in batches.
- 1 58. A method according to claim 52, further comprising:  
2 sending session information for that client session from the primary  
3 server to the secondary server in batches.
- 1 59. A method according to claim 52, further comprising:  
2 associating a version number with each update to the session  
3 information in response to a request.
- 1 60. A method according to claim 52, wherein sending session  
2 information for that client session from the primary server to the secondary  
3 server comprises sending a delta of information containing the changes in  
4 the session information.
- 1 61. A method according to claim 52, further comprising:  
2 assigning the primary and secondary servers each an identification  
3 number.
- 1 62. A method according to claim 52, further comprising:  
2 adding an identification number for the primary server and an  
3 identification number for the secondary server to obtain one  
4 number that describes both the primary and secondary  
5 servers.

- 1      63. A system for replicating information in a client session, comprising:  
2      a. a plurality of servers;  
3      b. a primary server in said plurality of servers, said primary  
4      server adapted to receive a request in a client session and  
5      serve a response to that request, said primary server further  
6      adapted to store session information for that client session,  
7      and adapted to generate a cookie containing information  
8      relating to said primary server and transmit the cookie to the  
9      session client;  
10     c. a secondary server in said plurality of servers selected by  
11     said primary server, said secondary server adapted to  
12     receive the session information from said primary server and  
13     store the session information, the secondary server further  
14     adapted to receive a request in a client session and serve a  
15     response to that request; and  
16     d. a hardware load balancer adapted to receive a request from  
17     a client, the request containing identification information for  
18     said primary and secondary servers, the web server adapted  
19     to process the identification information and serve the  
20     process request on said primary server, the hardware load  
21     balancer adapted to serve the request on a new primary  
22     server if said primary server is unable to process the request.
- 1      64. A method for providing redundancy in a client session, comprising:  
2      a. making a load balancing decision for an initial request from  
3      a client in a client session in order to select a primary server  
4      from a plurality of servers, the load balancing decision being  
5      made using an algorithm in a hardware load balancer;  
6      b. serving said request on said primary server;

- 7           c.     sending session information for that client session from the
- 8                 primary server to a secondary server selected by said
- 9                 primary server;
- 10          d.     storing a cookie on the session client containing information
- 11                 identifying the primary and secondary servers; and
- 12          e.     updating session information on the primary and secondary
- 13                 servers each time a request is received in the client session.

1     65.    A method according to claim 64, further comprising the step of  
2     processing a segment of the cookie containing identification information for  
3     a request received in the client session and serving the request on the  
4     primary server.

- 1     66.    A system for replicating information in a client session, comprising:
- 2           a.     a plurality of servers;
  - 3           b.     a primary server in said plurality of servers, said primary
  - 4                 server adapted to receive a request in a client session and
  - 5                 serve a response to that request, said primary server further
  - 6                 adapted to store session information for that client session,
  - 7                 and adapted to generate a cookie containing information
  - 8                 relating to said primary server and transmit the cookie to the
  - 9                 session client;
  - 10          c.     a secondary server in said plurality of servers selected by
  - 11                 said primary server, said secondary server adapted to
  - 12                 receive the session information from said primary server and
  - 13                 store the session information, the secondary server further
  - 14                 adapted to receive a request in a client session and serve a
  - 15                 response to that request; and
  - 16          d.     a hardware load balancer containing load-balancing logic for
  - 17                 selecting said primary server, said hardware load balancer
  - 18                 adapted to receive an initial request from a client and direct

19 that request to said primary server, said hardware load  
20 balancer further adapted to receive a subsequent request  
21 from a client, the subsequent request containing  
22 identification information for said primary and secondary  
23 servers, the hardware load balancer adapted to process the  
24 identification information and serve the process request on  
25 said primary server, the hardware load balancer further  
26 adapted to select a new primary server and serve the  
27 request on the new primary server if said primary server is  
28 unable to process the request.

- 1 67. A method for providing redundancy in a client session, comprising:
- 2 a. making a load balancing decision for an initial request from
- 3 a client in a client session in order to select a primary server
- 4 from a plurality of servers, the load balancing decision being
- 5 made using an algorithm in a hardware load balancer;
- 6 b. serving the request on the primary server;
- 7 c. sending session information for that client session from the
- 8 primary server to a secondary server selected by said
- 9 primary server;
- 10 d. storing a cookie on the session client, the cookie containing
- 11 information identifying the primary and secondary servers;
- 12 e. reading the identification information received with any
- 13 subsequent request in the client session and serving that
- 14 subsequent request on the primary server; and
- 15 f. updating session information on the primary and secondary
- 16 servers each time a subsequent request is served on the
- 17 primary server for that client session.

- 1 68. A method for providing redundancy in a client session, comprising:

- 2           a.     selecting a primary server from a plurality of servers in
- 3                 response to an initial request from a client, the selecting step
- 4                 using load-balancing logic in a hardware load balancer;
- 5           b.     serving said request on said primary server to begin a client
- 6                 session;
- 7           c.     selecting a secondary server;
- 8           d.     sending session information for that client session from the
- 9                 primary server to the secondary server; and
- 10          e.     updating session information on the primary and secondary
- 11                 servers each time a request is received in the client session.

1       69.    A method according to claim 68, further comprising:  
2               storing the information in a cookie on the client.


1       70.    A method according to claim 68, further comprising:  
2               serving the request on a new primary server if the primary server is  
3               unable to process the request.


1       71.    A method according to claim 68, further comprising:  
2               selecting a new primary server using load-balancing logic in the  
3               hardware load balancer.

1       72.    A method according to claim 68, further comprising:  
2               selecting a new secondary server.

1       73.    A method according to claim 68, further comprising:  
2               serving requests to the primary server in batches.

1       74.    A method according to claim 68, further comprising:  
2               sending session information for that client session from the primary  
3               server to the secondary server in batches.

*Figure 1*



*Figure 2*






Figure 3




Figure 4



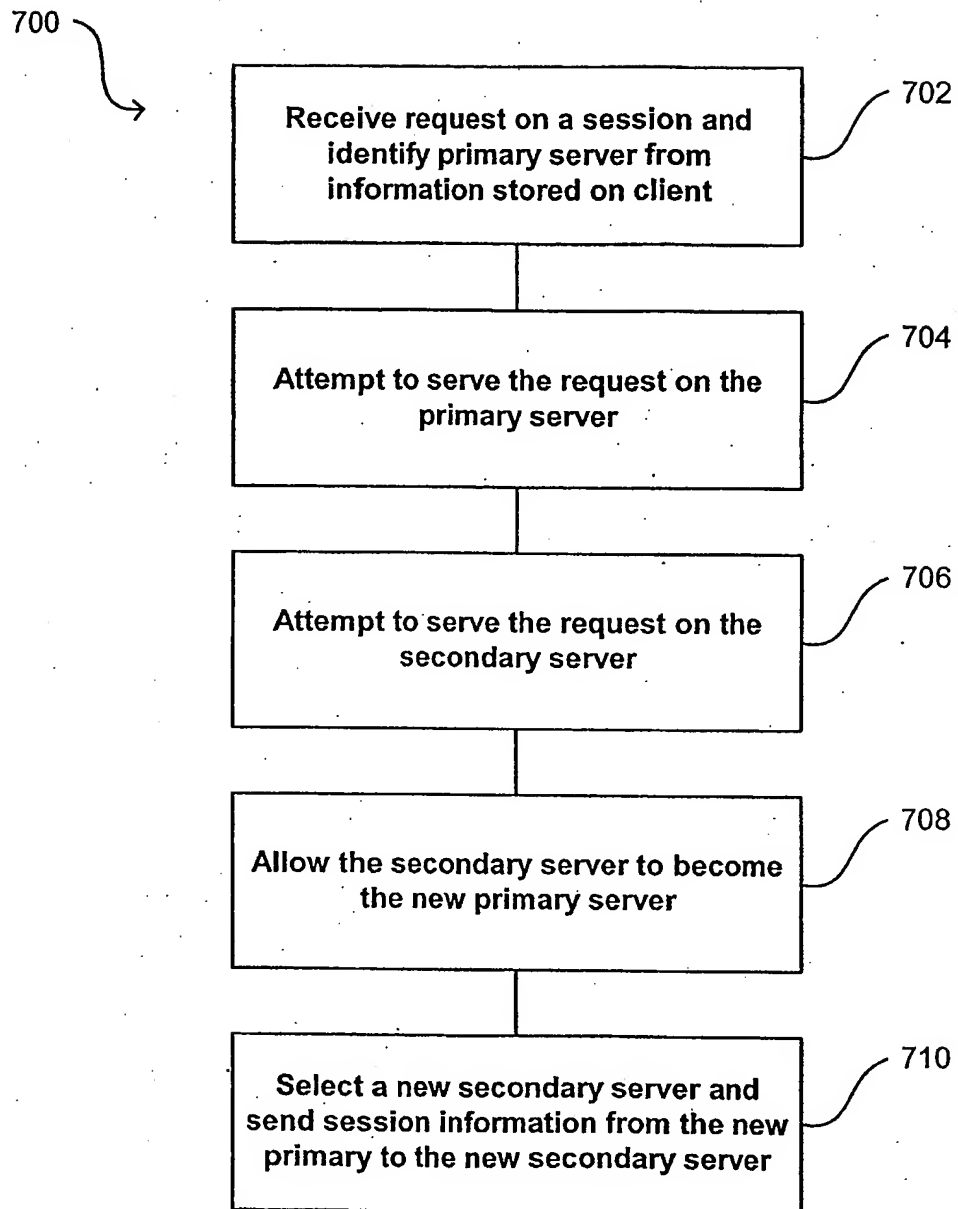
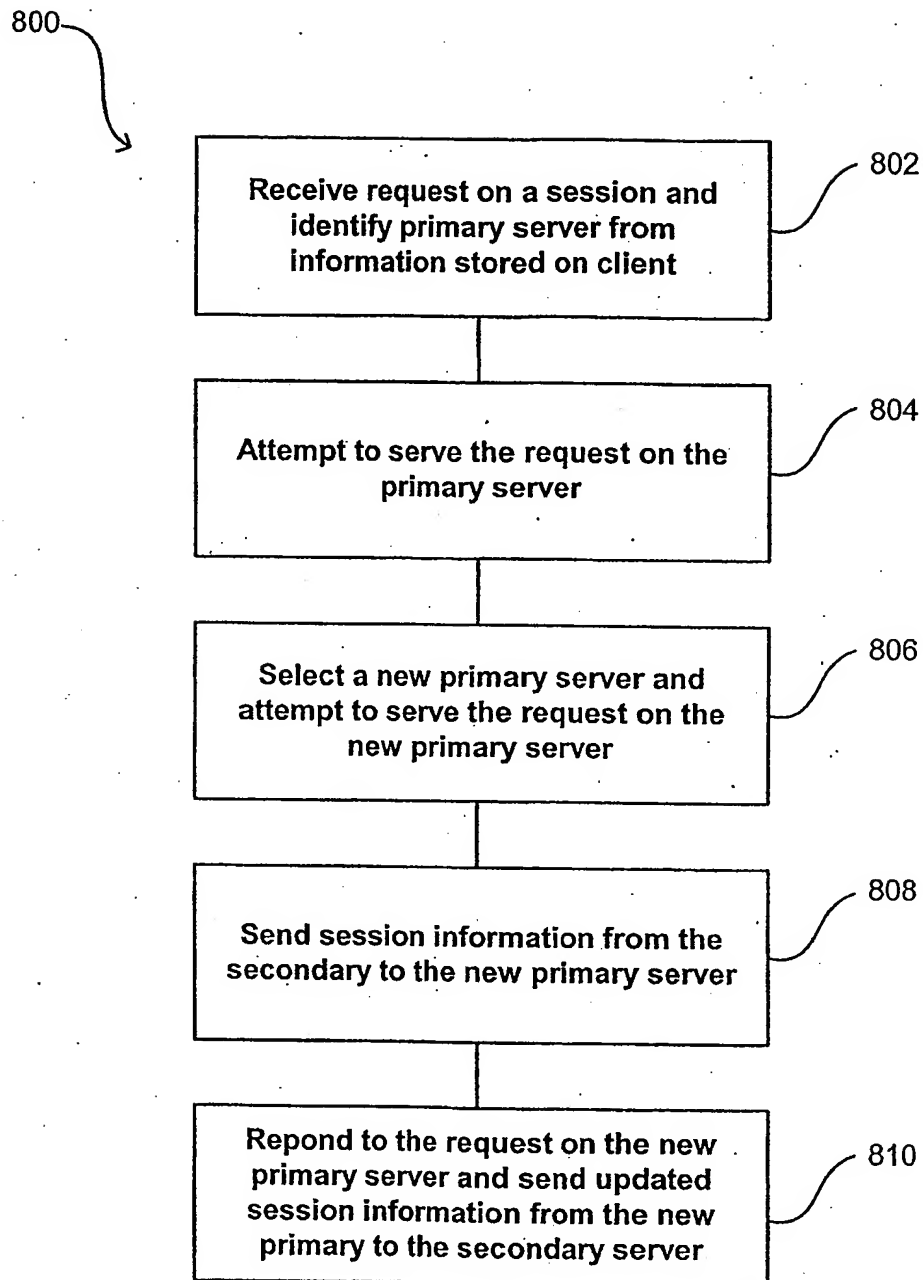


Figure 5

*Figure 6*

*Figure 7*

*Figure 8*

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/22429

**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(7) : G06F 15/16, 15/173

US CL : 709/105,226

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/105,226

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,951,694 A (CHOQUIER et al) 14 September 1999 (14.09.1999), see the whole reference.	1-74
A	US 5,828,847 A (GEHR et al) 27 October 1998 (27.10.1998), see the whole reference.	1-74
A	US 6,199,110 B1 (RIZVI et al) 06 March 2001 (06.03.2001), see the whole reference.	1-74
A,P	US 6,301,589 B1 (HIRASHIMA et al) 09 October 2001 (09.10.2001), see the reference.	1-74

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

27 September 2002 (27.09.2002)

Date of mailing of the international search report

12 NOV 2002

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks  
Box PCT  
Washington, D.C. 20231

Facsimile No. (703)305-3230

Authorized officer

Meng Ai An

Telephone No. 703-305-3900